

On Regularized Reconstruction of Vector Fields

Pouya Dehghani Tafti, *Member, IEEE*, and Michael Unser, *Fellow, IEEE*

Abstract—In this paper, we give a general characterization of regularization functionals for vector field reconstruction, based on the requirement that the said functionals satisfy certain geometric invariance properties with respect to transformations of the coordinate system. In preparation for our general result, we also address some commonalities of invariant regularization in scalar and vector settings, and give a complete account of invariant regularization for scalar fields, before focusing on their main points of difference, which lead to a distinct class of regularization operators in the vector case. Finally, as an illustration of potential, we formulate and compare quadratic (L_2) and total-variation-type (L_1) regularized denoising of vector fields in the proposed framework.

Index Terms—Curl and divergence in higher dimensions, fractional Laplacian, fractional vector calculus, regularization, rotation invariance, scale invariance, total variation (TV), vector fields, vector L_p spaces.

I. INTRODUCTION

OUR AIM in this paper is to derive, in a principled manner, formulas for regularization functionals suitable for reconstructing *vector fields*, with a view to applications such as denoising, deconvolution, and reconstruction from incomplete (that is, scalar) measurements [1], [2], among others. Our motivation in approaching the question of vector field reconstruction derives from the increasing prevalence of imaging modalities that produce measurements of vector quantities and the need to design algorithms for treating such data [3]. Such algorithms can also be applicable in other contexts where vector fields appear, such as estimating optical flow and image registration [4]–[7].

Throughout this paper, we take *invariance under coordinate transformations* as our guiding principle. The importance of invariance in reconstruction was already apparent to Duchon [8], who considered the problem of interpolating or approximating scalar fields in \mathbb{R}^d ; however, the mathematical formulation of invariance laws is, in general, different for scalars and vectors, as we shall see briefly in Section II and in more detail in Sections IV and V. The appeal of the notion of invariance partly lies in the fact that invariant regularizers do not impose a preferential choice of coordinate system on the model. We give a

rather complete characterization of invariant vector regularization operators in Sections IV and V, after initially reviewing the related scalar theory in Section III.

Regularized reconstruction of vector fields has been previously considered, notably by Suter and Chen [9], who proposed quadratic (L_2) regularization with mixed-order differentials of the vector field. Arigovindan *et al.* [2], [10] studied quadratic regularization with fractional-order differential operators and paid particular attention to the invariance properties of the regularization term with respect to vector rotation, translation, and change of scale, characterizing the complete family of quadratic regularization functionals with the required invariances, which essentially extend Duchon's thin-plate splines [8] to the vector setting. Specialized examples of such functionals, involving curl and divergence regularization, had been considered earlier by Dodu and Rabut [11] and (for the problem of interpolation) by Amodei and Benbourhim [12] before them.

All of the previous schemes fall under the general heading of smoothing spline and spline interpolation methods. They thus exhibit similar advantages (efficient resolution by linear methods and connection with splines) and limitations (most notably, oversmoothing of discontinuities and edges which, e.g., occur naturally at fluid interfaces in fluid dynamical systems and at object boundaries in optical flow). In this connection, it has been observed in the scalar setting that schemes using L_1 regularization—in particular, total variation (TV) type methods—do a better job of preserving edges and discontinuities than their L_2 counterparts [13], [14]. The framework that we have adopted in this paper allows us to find natural vector equivalents of these nonquadratic methods. (On the algorithmic side, the nonquadratic problems that we formulate here can be solved using techniques similar to those employed in the scalar case (see for instance [15]), as we show by way of examples in Section VI.)

On the theoretical side, another common property of quadratic schemes is that, due to the association of quadratic functionals with inner products, they can all be reduced to regularization with *self-adjoint* differential operators (essentially fractional Laplacians and their extensions; see Section IV). This is in contrast to the general nonquadratic case considered here, where the factorization of these self-adjoint operators into skew-symmetric ones becomes relevant (see Section V).

Finally, we wish to point out that unlike at least some of the previous works that have been exclusively concerned with 2-D and/or 3-D vector fields, the approach that we have adopted in the present paper makes it possible to consider vector fields in any number of dimensions on the same footing. This is particularly apparent in our dimensionless formulation of fractional Laplacians in Sections III and IV, and of *curl*- and *divergence-like* operators in Section V.

Manuscript received September 07, 2010; revised February 28, 2011; accepted May 11, 2011. Date of publication June 09, 2011; date of current version October 19, 2011. This work was supported in part by the Swiss National Science Foundation under Grant 200020-121763 and the ERC under Grant ERC-2010-AdG Z67439-FUN-SP. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Rafael Molina.

The authors are with the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland (e-mail: pouya.tafti@epfl.ch; michael.unser@epfl.ch).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2011.2159230

A. Regularized Reconstruction

The standard scenario for regularized reconstruction is as follows. We are given a vector y of measurements or observations that are assumed to depend, in a known probabilistic fashion (or deterministically but with some measurement and/or modeling error), on the unknown entity f_{true} , which we wish to reconstruct. f_{true} will be, in our case, a function defined on some finite or infinite domain. We then define the regularized reconstruction of f_{true} as the (hopefully unique) minimizer over f of a cost functional

$$\mathfrak{J}_\lambda(f; y) := \mathfrak{d}(f; y) + \lambda \mathfrak{R}(f) \quad (1)$$

composed of a *fidelity criterion* \mathfrak{d} quantifying the proximity of the observed measurements y to hypothetical measurements made from some possible reconstruction f and a *regularization functional* \mathfrak{R} that measures the undesirability of f based on our (deterministic or probabilistic) prior information or assumptions about the solution. The above formulation can be arrived at in different ways, some of which we shall now mention in passing for the sake of motivation, while reminding the reader that our primary purpose here is to derive some specific families of regularization functionals and not to justify the regularized variational framework for reconstruction in general (for comparable classifications see [4] and [16]).

- 1) In finite sample/parameter dimensions, that is, when both f and y are finite vectors, it is often possible to view the minimization of (1) as a case of maximum *a posteriori* (MAP) estimation. In this interpretation, \mathfrak{d} and $\lambda \mathfrak{R}$ essentially play the respective roles of the negative log likelihood and the negative log of the prior, usually up to some normalization (and possibly also discarding some terms that do not modify the solution). To come up with a prior, one might seek an operator R that *whitens* the vector f (i.e., renders its components independent); the log of the joint probability of the elements of Rf then becomes additive due to independence. This nicely fits with the most common form of regularization functionals used in practice, i.e., sums of the form

$$\sum_i \Phi([Rf]_i) \quad (2)$$

where R is a *regularization operator*, $[Rf]_i$ is the i th element of Rf , and Φ is a potential function such as $|\cdot|^2$ for ℓ_2 regularization or the absolute value for ℓ_1 regularization. Although MAP estimation is not the only purely probabilistic interpretation of (1), it is by far the most common one and, hence, the only one we shall mention here.

- 2) The form given in (1) can also be justified from a hybrid probabilistic–deterministic standpoint, where \mathfrak{d} again represents a negative log likelihood, whereas \mathfrak{R} now corresponds to the constraint

$$\mathfrak{R}(f) \leq a \quad (3)$$

on the solution, put in a Lagrange form with λ serving as the Lagrange multiplier. Such schemes are known under

the names of constrained or penalized likelihood. In addition, a connection can often be made with Grenander’s method of sieves [17] (where one considers a limiting sequence of minimizers of the cost functional with varying λ). Note that the roles of the constraint and the objective may be inverted, e.g., when the noise variance is known.

- 3) Finally, a purely deterministic interpretation is also possible, where $\lambda \mathfrak{R}$ is again the Lagrange relaxation of constraint $\mathfrak{R}(f) \leq a$, whereas \mathfrak{d} is a deterministic measure of data fidelity such as the Euclidean distance between y and samples of f . However, we remark that in many practical situations, the constraint bound a on which λ depends is not known (or the constraint is not really a hard one); consequently, λ can also be seen as a tuning parameter of the reconstruction algorithm.

Among the above justifications for the regularized reconstruction framework, the MAP interpretation does not trivially generalize to the case where an infinite number of values need to be estimated, which occurs, for example, when the domain of f is an infinite set such as \mathbb{R}^d (rather than a finitely countable set); for one thing, it is generally not possible to associate a probability distribution function, in its finite-dimensional sense, with probabilities on the function space to which f belongs, due to the fact that the Lebesgue measure does not admit of an infinite-dimensional generalization.

It is therefore constructive, in what follows, to imagine that the term $\lambda \mathfrak{R}(f)$ is derived from an inequality constraint as in the second and third interpretations. Moreover, we shall consider all algorithms based on the same \mathfrak{R} at the same time and consider λ (or equivalently, the constraint bound) as a tuning parameter of the algorithm. Probabilistic considerations then become secondary to geometrical/analytical ones, for which reason they shall not be emphasized in the remainder of this paper.

Even so, we still draw inspiration from the observation made at the end of paragraph 1) to define our regularization functionals as integrals of the form

$$\mathfrak{R}(f) = \int_{\text{Dom}f} \Phi(Rf(x)) dx \quad (4)$$

where R (previously the whitening operator) is now referred to as the *regularization operator*. Formally, the above integral, which replaces the sum in (2), can be thought of as the normalized aggregate contribution of individual independent pointwise innovations, i.e., the values of $Rf(x)$ as a function of x (although, strictly speaking, without proper normalization, such a contribution should be infinite from the probabilistic point of view).

After this brief introduction to regularized reconstruction, let us now describe the direction and contents of this paper. Our primary focus in this paper is vector field regularization. Thus, assuming the general form given in (4) for the reason that we just described, our task is then to specify the linear operator R and the function Φ . We shall derive the general form of admissible R s and Φ s—for scalars as well as for vectors—by imposing invariances under certain geometric transformations, namely, rotation (and reflection), translation, and scaling.

The motivation behind using invariances is that, in many physical systems, there exists no obvious preferential choice of

direction, position, or scale, at least within a reasonably wide range relevant in many applications. We therefore seek reconstruction algorithms that lead to a consistent solution under such transformations, possibly by appropriately adjusting a single parameter (λ). This requires the regularization functional \mathfrak{R} to be invariant under such transformations (possibly up to a computable multiplicative factor).

Although our main goal here is to formulate regularization functionals for vector fields, we begin the exposition by general considerations that apply equally to scalars and vectors (see Section II) and, for completeness, proceed to include a detailed account of invariant scalar regularization in Section III, where we derive the general form of R and Φ for the scalar case under suitable assumptions. Next, in Section IV, we turn our attention to vector fields and invariances relevant for them. This is followed by some extensions of the framework in Section V, where we additionally consider regularization operators that map vector fields to scalars and tensors. It will become clear by the end of Section V that, with a high level of generality, the functional \mathfrak{R} takes the form $\|Rf\|_p^q$ where we have the following.

In the scalar setting, $\|\cdot\|_p$ is the standard Lebesgue L_p norm, and R is either a fractional scalar Laplacian (defined in Section III) or a fractional gradient (defined in Section V).

In the vector setting, $\|\cdot\|_p$ is a suitable generalization of the scalar L_p norm to vector- or matrix-valued functions (introduced, respectively, in Section II and Appendix A), and R is either a generalized fractional vector Laplacian (introduced in Section IV), which incorporates a Helmholtz decomposition into curl- and divergence-free components, or else, it is a fractional curl or a fractional divergence (both introduced in Section V).

We then illustrate the proposed construction in Section VI, where we consider the problem of vector field denoising in 2- and 3-D and compare two solutions (quadratic and TV-like) that fall within our framework. Some remarks in Section VII conclude this paper.

Symbols and other notation are defined when first used and summarized in Table I for reference.

II. GENERALITIES REGARDING REGULARIZATION AND INVARIANCE

As noted in the Introduction, in identifying suitable families of regularization functionals, we are guided by the principle of invariance under specific geometric transformations. With any such transformation is associated a symbol S that can be a scalar $\sigma > 0$ (the scale) for changes of scale, a vector $\tau \in \mathbb{R}^d$ (the displacement vector) in the case of translations, an orthogonal transformation matrix $\omega \in \mathbb{R}^{d \times d}$ when considering rotations and reflections, or, once again, a scalar $\alpha > 0$ (the gain) when multiplication by positive reals (change of units) is considered.

Since, in general, the *same* transformation group can act differently on scalars and vectors (this is particularly true for rotations, as we shall see in Section IV), the same symbol S can describe different laws of transformation, depending on whether it is acting on scalars or vectors or other entities. For this reason, we introduce the notation $[S]$ to denote the *operator* associated

TABLE I
NOTATION

symbol	description
d	number of spatial dimensions
\mathbb{R}_+	set of positive reals = $\{a \in \mathbb{R} a > 0\}$
$x = (x_1, \dots, x_d)$	spatial coordinates
$\xi = (\xi_1, \dots, \xi_d)$	Fourier coordinates (dual to x)
$f = f(x)$	field of scalars, vectors, or bivectors (usu. denoting a possible reconstruction)
$\hat{f} = \hat{f}(\xi)$	Fourier transform of f
$\Omega = \{\omega\}$	group of orthogonal matrices $\omega \in \mathbb{R}^{d \times d}$
$T = \{\tau\}$	group of displacement vectors $\tau \in \mathbb{R}^d$
$\Sigma = \{\sigma\}$	group of spatial scale factors $\sigma \in \mathbb{R}_+$
$A = \{\alpha\}$	group of amplitude gain factors $\alpha \in \mathbb{R}_+$
S	placeholder for ω , τ , σ , or α
$[S]f$	transformation of f by S ; $[S]$ can be understood as the <i>operator</i> that transforms f
$[S]_s f$	same as above, additionally indicating that f is scalar-valued
$[S]_v f$	same as above, additionally indicating that f is vector-valued
$[S]_b f$	same as above, additionally indicating that f is bivector-valued
$ a $	for $a \in \mathbb{R}$, the absolute value of a for $a \in \mathbb{R}^d$, the Euclidean length of a
$\ f\ _p$	for scalar-valued f , the standard Lebesgue L_p norm of f for vector-valued f , the L_p norm defined in Corollary 1 for matrix-valued f , the L_p norm defined in Appendix A
$\langle a, b \rangle$	the scalar product of vectors $a, b \in \mathbb{R}^d$
$\langle f, g \rangle$	the scalar product of functions f, g (= $\int_{\mathbb{R}^d} \overline{f^T} g$)
R	regularization operator
div	divergence operator (vector to scalar)
grad	gradient operator (scalar to vector)
curl	curl operator (vector to bivector)
curl*	adjoint curl operator (bivector to vector)
*	star operation (bivector to pseudo-vector, see (21))

with symbol S and distinguish between scalar and vector operators by using subscripts as per $[S]_s$ and $[S]_v$ where necessary.

Definition 1: In mathematical terms, we assume that S belongs to one of several *transformation groups* $T = \{\tau \in \mathbb{R}^d\}$ (the *translation group*), $\Omega = \{\omega \in \mathbb{R}^{d \times d} : \omega \text{ orthogonal}\}$ (the *orthogonal group*), $\Sigma = \{\sigma \in \mathbb{R}_+\}$ (the *scaling group*), or $A = \{\alpha \in \mathbb{R}_+\}$ (the *gain group*), and consider maps (isomorphisms) $S \mapsto [S]_C$ between transformation groups and groups of operators (*actions* or *transformation laws*) acting on objects of some class C ($C = s$ for scalar fields, $= v$ for vector fields, $= b$ for bivector fields, etc.).

We then define $[\tau]f(\cdot) = f(\cdot - \tau)$ for $\tau \in T$ (translation); $[\alpha]f(\cdot) = \alpha f(\cdot)$ for $\alpha \in A$ (gain); $[\sigma]f(\cdot) = f(\sigma^{-1} \cdot)$ for

$\sigma \in \Sigma$ (scaling); and $[\omega]_s f(\cdot) = f(\omega^T \cdot)$ for $\omega \in \Omega$ (orthogonal transformation of scalars). Note that the first three identities are valid for scalars and for vector fields alike, whereas the last one only applies to scalars. Vector rotation follows a different rule: $[\omega]_v f(\cdot) = \omega f(\omega^T \cdot)$. The reason is that the coordinates f_1, f_2, \dots, f_d of a vector field $f(x)$, $x \in \mathbb{R}^d$, are specified in the same coordinate system as that of its argument x , which means that if the coordinate system of the argument is rotated by ω^T , the coordinates f_1, \dots, f_d have to be transformed by the inverse (ω) in order to keep the direction of the vectors fixed. \square

We recall [cf. (4)] that we shall be seeking invariant regularization functionals of the following form:

$$\mathfrak{R}(f) = \int_{\mathbb{R}^d} \Phi(Rf(x)) dx$$

where the scalar-valued function Φ and the operator R are to be determined.

Requiring that the regularization be S -invariant up to some readjustment of parameter λ amounts to demanding that

$$\mathfrak{R}([S]f) = c_{S,\mathfrak{R}} \mathfrak{R}(f) \quad (5)$$

for all f under consideration, where $c_{S,\mathfrak{R}}$ is a constant. In order to have more flexibility in constructing regularization functionals, we wish to find families of functions Φ and operators R that we can then pick and combine independently. In particular, since we shall always include *identity* in our family of regularization operators, we require Φ to satisfy the following:

$$\int_{\mathbb{R}^d} \Phi([S]f(x)) dx = c_S \int_{\mathbb{R}^d} \Phi(f(x)) dx \quad (6)$$

for all f and all S and for some constant $c_S > 0$ that depends on S .

From (6) immediately follows Proposition 1.

Proposition 1: Let Φ satisfy (6) and be continuous on some open neighborhood. Φ is then equivalent to a homogeneous function; that is, $\Phi(a) = c|a|^p$ (almost everywhere) for some $c > 0$ and $p \in \mathbb{R}$ ($|a|$ denotes the absolute value or the modulus of a as appropriate).

Conversely, (6) holds for any such Φ as long as the integrals are well-defined. \square

Proof: From (6), we have $\int [\Phi([S]f(\cdot)) - c_S \Phi(f(\cdot))] = 0$ for all f , and therefore

$$\Phi([S]f(x)) = c_S \Phi(f(x)) \quad \text{for almost all } x.$$

We shall first consider the case of scalar f , where Φ is a function of the reals. Let $S = \alpha$ belong to the gain group A , and let $\phi(\alpha) := c_\alpha > 0$. We then have, for arbitrary $\alpha = -1$ and arbitrary $a = f(x)$, $\Phi(a) = \Phi(-(-a)) = [\phi(-1)]^2 \Phi(a)$ whence $\phi(-1) = 1$ and $\Phi(a) = \Phi(|a|)$. Next, for arbitrary α and a , we may write the following:

$$\phi(|\alpha|)\Phi(|a|) = \Phi(\alpha a) = \Phi(a\alpha) = \phi(|a|)\Phi(|\alpha|).$$

Fixing either α or a then proves that $\phi = c'\Phi$ for some constant c' . Therefore, for all $\alpha, a \in \mathbb{R}_+$

$$\Phi(\alpha a) = c'\Phi(\alpha)\Phi(a).$$

This shows that Φ is an exponential function and can therefore be written as $\Phi(\cdot) = c(\cdot)^p$ for some constants c and p , as claimed.

When f is vector-valued, rotation invariance implies that Φ is in fact only a function of the modulus $|\alpha|$ of α ; we may then repeat the argument of the previous paragraph to once again deduce that $\Phi(\alpha) = c|\alpha|^p$ for some c and p .

To prove the converse, one can directly inspect each of the groups of transformations involved by a simple change of variables in the integrals and verify that the desired result follows from the invariances of the Lebesgue measure. \blacksquare

The following corollary is immediate.

Corollary: Vector norms

$$\|f\|_p := \begin{cases} \left(\int_{\mathbb{R}^d} |f(x)|^p dx \right)^{1/p}, & 1 \leq p < \infty, \\ \text{ess sup } |f(x)|, & p = \infty, \end{cases} \quad (7)$$

are S -invariant in the sense that $\|[S]f\|_p = c_{S,p}\|f\|_p$ for all (vector-valued) f . Conversely, any convex S -invariant integral functional [as defined in (6)] that satisfies the requirements of Proposition 1 is of the form $\|f\|_p^q$ for some $p \geq 1$. \square

It is then sufficient, in order to have the desired independence between the choice of Φ and R , to require that R commute with coordinate transformations up to a multiplicative constant $k_S \neq 0$, in the sense that

$$R[S]f = k_S[S]Rf \quad (8)$$

for all f . This, we note, is the quintessence of invariance, as it means that applying the coordinate transformation before or after the application of R yields the same result (up to normalization).

Consequently, the regularization functional given in (4) can be written as (the p th power of) the L_p norm of Rf (we absorb all the constants in λ ; p is required to be ≥ 1 for the sake of convexity). We may also include the ∞ -norm $\mathfrak{R}(f) = \|Rf\|_\infty$ for completeness since, even though it is not strictly derived from an integral, it nevertheless satisfies the required invariances.

It is worth noting that, following the Lagrangian interpretation given in the introduction [cf. (3)], we may, in practice, replace $\mathfrak{R}(f)$ by $\Psi(\mathfrak{R}(f))$, where Ψ is an arbitrary continuous strictly increasing function on \mathbb{R}_+ , since all such functions define equivalent inequality constraints in the Lagrangian formulation, for $\mathfrak{R}(f) \leq a \Leftrightarrow \Psi(\mathfrak{R}(f)) \leq \Psi(a)$. Such a function Ψ can therefore be introduced as convenient. However, if it is desired to have (5) hold with a constant $c_{s,\mathfrak{R}}$ not depending on f , one can then show that Ψ needs to be a multiple of the homogeneous function $|a|^q$ for some q (cf. the proof of Proposition 1). Putting all this together, we get the following.

Proposition 2: Let R be S -invariant in the sense of (8). Then, given $p \in [1, \infty]$ and any $q \in \mathbb{R}_+$, the regularization functionals

$$\mathfrak{R}_p^q(f) := \|Rf\|_p^q \quad (9)$$

are S -invariant up to a multiplicative factor; that is, we have

$$\mathfrak{R}_p^q([S]f) = c_{S,p,q} \mathfrak{R}_p^q(f)$$

for some $c_{S,p,q} \neq 0$. \square

Proof: This is an immediate consequence of Corollary 1 and (8). ■

Excepting the case of $p = \infty$, where one normally takes $q = 1$, the preferred choice of q in practice is $q = p$, which simplifies the formulas by getting rid of the p th algebraic root hidden in the definition of the L_p norm.

As a reminder, in (8) (reproduced below for convenience), we required that the operator $R : X \rightarrow Y$ commute with the transformation associated with S , where S is taken from one of the transformation groups Σ, A, T , or Ω (cf. Definition 1), i.e.,

$$R[S]_X = k_S[S]_Y R. \tag{10}$$

Note that, in general, when R maps objects of type X to those of a different type Y (such as vectors to scalars or vice versa), the operator associated with S will be different on the two sides of (10); we have emphasized this in the above equation by subscripting the operator with X and Y as appropriate.

We say that R is $\{S\}$ invariant if it satisfies (10) for all S in some understood transformation group(s) (*strictly* $\{S\}$ invariant if, in addition, $k_S = 1$). For instance, we shall talk about Ω -invariant (T -invariant, etc.) operators, by which we mean operators that satisfy (10) for $S \in \Omega$ ($S \in T$, etc.). One notes that, for an $\{S\}$ -invariant operator, the map

$$S \mapsto k_S \tag{11}$$

is a group homomorphism from any of the transformation groups under consideration (typically, Σ, A, T , or Ω) onto (a subset of) \mathbb{R} .

In the sequel, we shall limit ourselves to *linear* regularization operators R while reminding the reader that, in general, the reconstruction problem remains nonlinear due to the L_p norms involved. We shall also assume that R is stable under shifts in the sense defined below.

Definition 2: Operator R is said to be *minimally T -stable* in L_p if there exists a subset E of L_p , not entirely inside the kernel of R , that is invariant under the action of T and on which R has a bounded operator norm; that is, if the following conditions are simultaneously satisfied:

$$\begin{aligned} [\tau]f \in E & \quad \text{for all } f \in E \text{ and all } \tau \in T; \\ \frac{\|Rf\|_p}{\|f\|_p} < C & \quad \text{for some } C < \infty \text{ and all } f \in E; \\ 0 < \|Rf\|_p & \quad \text{for some } f \in E. \end{aligned} \tag{12}$$

□

In some problems of practical interest, one may wish to consider a combination of N regularization terms, rather than a single one of them. These different regularizers may, for instance, measure the regularity of the projections of f onto different subspaces with special physical significance (we shall see some examples of these in Section IV, where we consider curl- and divergence-free subspaces). In this case, the cost functional to be minimized takes the following form:

$$\mathfrak{J}_{(\lambda_1:\lambda_N)}(f; y) = \mathfrak{d}(f; y) + \sum_{1 \leq i \leq N} \lambda_i \|R_i f\|_{p_i}^{q_i}$$

which can also be interpreted as the Lagrange relaxation of a constrained optimization problem with several inequality constraints (i.e., $\|R_i f\|_{p_i}^{q_i} \leq a_i, 1 \leq i \leq N$). Since, per Proposition 2, each of the regularization terms is invariant under the desired geometric transformations, their weighted sum will also have this property, up to a suitable independent adjustment of the λ_i 's for each given geometric transformation. As such, all that was or will be said here, in connection with the interplay of invariance and regularization, will be understood to generalize in the sense just described to linear combinations of regularization terms.

Having established the general form of regularization functionals in terms of L_p norms of Rf , where R is the regularization operator with invariance properties dictated by (10), we shall now take up the task of identifying such operators. This will require us to consider scalar and vector cases separately, primarily due to the difference in the law of rotation in the two settings.

III. REGULARIZATION OPERATORS: SCALAR CASE

Here, we shall derive the general form of linear regularization operators that possess specific invariance properties in the sense of (10). Our main result here is stated in Theorem 1, which shows that these operators take the form of fractional Laplacians.

We refer the reader to Definition 1 for a list of invariances that are of interest to us. Some peculiarities of the translation group T and the orthogonal group Ω , together with the stability assumption described in Definition 2, allow us to show in the following that, for transformations in these two groups, the constant k_S in (10) is always 1.

Lemma 1: A minimally T -stable operator R (cf. Definition 2) that is invariant under the action of Ω and T in the sense of (10) is *strictly* invariant under Ω and T , that is, it has

$$k_S = 1$$

for all $S \in \Omega \cup T$. □

Proof: First, note that for those elements of Ω that are of some finite order m , i.e., for any orthogonal matrix ω such that $\omega^m = \text{Id}$, by the $\omega \mapsto k_\omega$ homomorphism [cf. (11)], we have $1 = k_{\text{Id}} = k_{\omega^m} = k_\omega^m \Rightarrow k_\omega = 1$ (knowing that $k_\omega \in \mathbb{R}_+$).

Furthermore, any element ω of Ω , including those of infinite order, can be written as a product of, at most, d reflections ω_i , where d is the dimension (this is the Cartan–Dieudonné theorem). Reflections are of order 2 and hence have coefficient $k_{\omega_i} = 1$ by the previous paragraph. We therefore have, for arbitrary $\omega \in \Omega$, $k_\omega = \prod_i k_{\omega_i} = 1$. This proves the Ω part of the lemma.

We shall prove the second part by contradiction. To this end, assume that there exists $\tau \in T$ with $k_\tau \neq 1$. Without loss of generality, we may assume $k_\tau > 1$ (simply replace τ by $-\tau$ in the other case). Then, for some $f \in E$ not in the kernel of R , with E defined in Definition 2, we have

$$\lim_{m \rightarrow \infty} \frac{\|R[m\tau]f\|}{\|[m\tau]f\|} = \lim_{m \rightarrow \infty} k_\tau^m \frac{\|Rf\|}{\|f\|} \rightarrow \infty$$

which contradicts (12).

Finally, note that, if we had restricted ourselves to rotation matrices instead of general orthogonal transformations in the first part of the lemma, we could still have proved $k_\omega = 1$ with the aid of an additional minimal Ω -stability assumption, arguing as we did for T . ■

We also have the following.

Lemma 2: The factor k_σ corresponding to scaling with $\sigma > 0$ [cf. (10)] is homogeneous in σ , that is, it can be written as

$$k_\sigma = \sigma^{-\gamma}$$

for some $\gamma \in \mathbb{R}$. □

The proof is very similar to that of Proposition 1; hence, we omit it.

The stage is now set for the following result. In some form, this result goes back to Duchon, although here we derive it from somewhat different premises (such as minimal T -stability).

Theorem 1: Let R be a real and minimally T -stable Fourier integral operator, initially defined from the Schwartz space \mathcal{S} to L_p for some $p \geq 1$, which is invariant under the action of T , Ω , and Σ in the sense of (10). R is then characterized by a Fourier multiplier of the following form:

$$\hat{R}(\xi) = c|\xi|^\gamma \quad (13)$$

where $\gamma > d - d/p$ is the exponent identified in Lemma 2.

Conversely, Fourier operators with symbols given by (13) are strictly invariant under the action of T and Ω and invariant under the action of Σ with the same coefficient k_σ as in Lemma 2. □

Proof: First, observe that, by Lemma 1, R is strictly T - and Ω -invariant, and by Lemma 2, its Σ -invariance coefficient k_σ is a homogeneous function $\sigma^{-\gamma}$ of $\sigma \in \mathbb{R}_+$. Since R is a linear and translation-invariant Fourier operator, it is associated with an integral as per

$$\begin{aligned} Rf(x) &= (2\pi)^{-d} \int_{\mathbb{R}^d} e^{i\langle x, \xi \rangle} \hat{R}(\xi) \hat{f}(\xi) d\xi \\ &= (2\pi)^{-d} \left\langle \hat{R}, e^{i\langle x, \cdot \rangle} \hat{f} \right\rangle \end{aligned}$$

where \hat{R} is the Fourier multiplier corresponding to R .

One can then directly verify that in order for R to commute with rotations and scalings (the latter up to a homogeneous multiplicative factor of $k_\sigma = \sigma^{-\gamma}$), its Fourier expression \hat{R} must be rotationally symmetric and homogeneous of degree γ . It is known [18], [19] that, subject to L_p boundedness, all such distributions can be represented in the following form:

$$\hat{R}(\xi) = c|\xi|^\gamma \quad \text{with } \gamma > d - d/p.$$

The same proof goes through when restricting ourselves to rotations instead of general orthogonal transformations if we make the additional assumption of minimal Ω -stability.

The converse is easily verified by simple changes of variables in the Fourier domain. ■

Note that $|\xi|^\gamma$ is the Fourier symbol of the $(1/2)\gamma$ th (fractional) power of the negative Laplacian $(-\Delta)$. We can therefore write the reconstruction cost functional as

$$\mathfrak{J}_\lambda(f; y) = \mathfrak{d}(f; y) + \lambda \left\| (-\Delta)^{\gamma/2} f \right\|_p^q.$$

Moreover, by the argument given at the end of the previous section, we may additionally consider multiple additive regularization terms, as in

$$\mathfrak{J}_{(\lambda_1: \lambda_N)}(f; y) = \mathfrak{d}(f; y) + \sum_{1 \leq i \leq N} \lambda_i \left\| (-\Delta)^{\gamma_i/2} f \right\|_{p_i}^{q_i}.$$

Two of the most important regularization functionals traditionally used in image processing are the total variation of f and its L_2 counterpart: $\|\text{grad} f\|_1$ and $\|\text{grad} f\|_2^2$ [using the vector L_p norms of (7)], both of which satisfy the required invariances. Note, however, that these regularizers, as such, fall outside the scope of this section for the reason that they incorporate an operator (grad) that maps scalars to vectors, whereas all the operators considered so far map to scalars and not vectors. Nevertheless, due to a peculiar property of the L_2 norm (namely, that it is a Hilbert space and has an inner product structure), in the L_2 case, one can write the following:

$$\begin{aligned} \|Rf\|_2^2 &= \langle Rf, Rf \rangle = \langle R^* Rf, f \rangle \\ &= \left\langle (R^* R)^{\frac{1}{2}} f, (R^* R)^{\frac{1}{2}} f \right\rangle = \left\| (R^* R)^{\frac{1}{2}} f \right\|_2^2 \quad (14) \end{aligned}$$

where R^* is the adjoint of R , and the self-adjoint operator $(R^* R)^{1/2}$ maps scalars to scalars and is therefore included in our framework. Hence, in particular, for the L_2 grad regularizer, we have $\|\text{grad} f\|_2^2 = \|(-\Delta)^{1/2} f\|_2^2$, which belongs to the family we derived above; the same cannot be said about L_1 TV.

Partly in order to overcome the latter limitation, later, in Section V, we shall also develop the theory of scalar-to-vector regularization operators and introduce fractional gradients grad^γ . From there, it then follows that, more generally, invariant scalar cost functionals can be of the form

$$\begin{aligned} \mathfrak{J}_{(\lambda_1: \lambda_N)}(f; y) &= \mathfrak{d}(f; y) + \sum_{1 \leq i \leq N} \lambda_i \left\| (-\Delta)^{\gamma_i/2} f \right\|_{p_i}^{q_i} \\ &\quad + \sum_{1 \leq i \leq N'} \lambda'_i \left\| \text{grad}^{\gamma'_i} f \right\|_{p'_i}^{q'_i}. \end{aligned}$$

IV. REGULARIZATION OPERATORS: VECTOR CASE

Translations $\tau \in T$ and scalings $\sigma \in \Sigma$ act in the same way on vector fields as they do on scalars. On the other hand, we shall need to redefine the action of orthogonal group Ω in the vector setting. Since a vector field is specified in the same coordinate system in which its argument is given, when transforming the domain, one has to recompute the coordinates of the vector

field accordingly. More precisely, the formula for transforming vector field $f = (f_1, \dots, f_d)^T$ by orthogonal matrix ω is

$$[\omega]_v f = \omega f(\omega^T \cdot) \quad (15)$$

that is, the coordinates of the vector are transformed by the inverse of the domain transformation matrix. On occasion, we shall refer to invariance as in (15) as *contra-variance* (recall that we distinguish between the scalar and vector operators associated with ω by subscripting $[\omega]$ by s and v , respectively).

The following result, indirectly proved for $d = 2, 3$ in [10], is the vector counterpart of Theorem 1. In the Appendix, we give a different and more general proof of this theorem, which is valid in any number of dimensions.

Theorem 2: Let R be a real and minimally T -stable Fourier operator initially defined $\mathcal{S}^d \rightarrow L_p^d$ and mapping vector fields to vector fields, which is invariant under the action of T , Ω , and Σ in the sense of (10). R is then characterized by a (matrix-valued) Fourier multiplier of the following form:

$$\hat{R}(\xi) = |\xi|^\gamma \left[c_1 \frac{\xi \xi^T}{|\xi|^2} + c_2 \left(I - \frac{\xi \xi^T}{|\xi|^2} \right) \right] \quad (16)$$

where c_1 and $c_2 \in \mathbb{R}$ are Helmholtz coefficients (see below) and $\gamma > d - d/p$ is the exponent identified in Lemma 2.

Conversely, operators with Fourier multipliers as above satisfy all of the required invariances. \square

Sketch of the proof: The complete proof appears in the Appendix. Here is an introduction to it.

The $k_\tau = 1$ part of Lemma 1 and Lemma 2 (which says that $k_\sigma = \sigma^{-\gamma}$ for some γ) apply without modification in the vector setting. Furthermore, following the same line of argument as in the proof of Lemma 1, one can prove that once again, $k_\omega = 1$ for all $\omega \in \Omega$, as was the case for scalars. Since R is linear and translation invariant, it admits a Fourier-domain representation as

$$Rf(x) = (2\pi)^{-d} \int_{\mathbb{R}^d} e^{i\langle x, \xi \rangle} \hat{R}(\xi) \hat{f}(\xi) d\xi$$

where \hat{R} is now a *matrix-valued* Fourier kernel. The scale invariance of R with coefficient $k_\sigma = \sigma^{-\gamma}$ and its strict orthogonal contra-variance translate to the following Fourier-domain identities:

$$\hat{R}(\sigma \cdot) = \sigma^{-\gamma} \hat{R}(\cdot) \quad \text{for all } \sigma > 0 \quad (17)$$

$$\hat{R}(\omega \cdot) = \omega \hat{R}(\cdot) \omega^T \quad \text{for all } \omega \in \Omega. \quad (18)$$

In the Appendix, we prove the forward direction in two steps, first showing that orthogonal contra-variance implies that $\hat{R}(\xi)$, at any $\xi \neq 0$, has an eigendecomposition as

$$\hat{R}(\xi) = \mu_1(\xi) \frac{\xi \xi^T}{|\xi|^2} + \mu_2(\xi) \left(I - \frac{\xi \xi^T}{|\xi|^2} \right) \quad (19)$$

and then noting that, by Theorem 1, μ_1 and μ_2 must be of the form $c_i |\xi|^\gamma$, $i = 1, 2$.

The converse of the theorem can easily be verified by Fourier-domain changes of variables. \blacksquare

With regard to parameters c_1 and c_2 , three cases are of particular interest, namely, those of $c_1 = c_2$, $c_1 = 0$, and $c_2 = 0$.

For $c_1 = c_2 = c$, the operator defined in Theorem 2 has the Fourier expression $c|\xi|^\gamma$ and therefore corresponds, up to normalization, to the fractional vector Laplacian $(-\Delta)^{\gamma/2}$, that is, the scalar Laplacian applied coordinatewise. For this reason, we shall refer to the family of operators identified by (19) as *generalized vector Laplacians*, with notation $(-\Delta)_{(c_1, c_2)}^{\gamma/2}$ [20].

To better understand the behavior of the operator when either c_1 or c_2 is zero, note that $(-\Delta)_{(c_1, c_2)}^{\gamma/2}$ can be decomposed as

$$[c_1(\text{Id} - \mathbb{P}) + c_2\mathbb{P}] (-\Delta)^{\gamma/2} \quad (20)$$

where operator \mathbb{P} is defined by its Fourier multiplier $\hat{\mathbb{P}} = \xi \xi^T / |\xi|^2$. It is straightforward to see that \mathbb{P} and its complement $\text{Id} - \mathbb{P}$ are projections and that they, in fact, project their argument onto its curl- and divergence-free components, respectively; in other words, taken together, they provide a Helmholtz decomposition of their argument.

To summarize, the operators identified in Theorem 2 effectively combine a fractional vector Laplacian with a reweighting of Helmholtz components. Moreover, one has the following:

$$(-\Delta)_{(c_1, c_2)}^{\gamma/2} (-\Delta)_{(c'_1, c'_2)}^{\gamma'/2} = (-\Delta)_{(c_1 c'_1, c_2 c'_2)}^{(\gamma + \gamma')/2}.$$

We can now give the general form of our cost functional for vector fields, as we did for scalar fields in Section II. Once again, we may consider linear combinations of some N regularization terms, which retain the same invariances as the individual terms, up to readjustment of $\lambda_1, \dots, \lambda_N$ as

$$\mathfrak{J}_{(\lambda_1; \lambda_N)}(f; y) = \mathfrak{D}(f; y) + \sum_{1 \leq i \leq N} \lambda_i \left\| (-\Delta)_{(c_{1,i}, c_{2,i})}^{\gamma_i/2} f \right\|_{p_i}^{q_i}.$$

However, the above family is still not complete, for reasons similar to those given at the end of Section III. This consideration forms the basis of our next section.

V. MORE ON L_p REGULARIZATION OF VECTOR FIELDS

A. Motivation

In our discussion in the preceding sections, we implicitly assumed that R mapped scalar or vector fields to similar objects and in the same number of dimensions. In other words, we considered the operator associated with S in (10) to be the same on the left and right sides. In this way, we overlooked some important possibilities for vector regularization operators, such as the divergence operator (mapping vector fields to fields of scalars) or the curl (mapping vector fields to pseudovector fields in 3-D; see below). Here, we shall remedy this by studying operators that generalize divergences and curls (and their adjoints), in the same way that the operators of the preceding sections generalized scalar and vector Laplacians.

The generalization to d dimensions of the divergence raises no difficulty. Indeed, the divergence of a vector field is defined in any number of dimensions by means of Fourier multiplier $i\xi^T$ (given in Cartesian coordinates). The divergence maps vector fields to scalar fields. Its adjoint is the negative gradient with Fourier multiplier $-i\xi$, which maps scalar fields to vector fields.

It is less obvious how the usual 3-D definition of the curl can be generalized to d dimensions. This difficulty is essentially rooted in the fact that the curl of a vector in 3-D is not a *true* vector: Per the right-hand rule of physics, the curl of a vector field transforms as an ordinary vector field under proper rotations, but it flips sign under *improper* rotations (those with determinant -1). For this reason, curl fields in 3-D are usually referred to as *pseudovector* fields.

It is in fact this notion of pseudovector that does not directly generalize to arbitrary d . For this reason, in higher dimensions, it is constructive to consider the curl operator as a map from vector fields to *bivector* fields ($d \times d$ matrix fields with specific transformation laws). We may identify bivectors with fields of $d \times d$ antisymmetric tensors [21]. These have $d(d-1)/2$ independent components, corresponding to the upper-diagonal elements of the tensor (only in 3-D is $d = d(d-1)/2$; hence, the difficulty in generalizing the customary definition of the curl and pseudovectors to $d > 3$).

In three dimensions, identification between pseudovectors and antisymmetric matrices (bivectors) can be made by the \star -map that we introduce as

$$\frac{1}{\sqrt{2}} \begin{bmatrix} 0 & -f_3 & f_2 \\ f_3 & 0 & -f_1 \\ -f_2 & f_1 & 0 \end{bmatrix} \mapsto \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix}. \quad (21)$$

The d -dimensional generalization of the curl, as a map from vector fields to bivector fields, is then given by the Fourier expression

$$(\text{curl}f)^\wedge = \frac{1}{\sqrt{2}}(i\xi \hat{f}^T - \hat{f} i\xi^T).$$

Defining the vector-gradient $\text{grad}f$ of the vector field f as the matrix

$$\text{grad}f = \begin{bmatrix} \partial_1 f_1 & \partial_1 f_2 & \partial_1 f_3 \\ \partial_2 f_1 & \partial_2 f_2 & \partial_2 f_3 \\ \partial_3 f_1 & \partial_3 f_2 & \partial_3 f_3 \end{bmatrix}$$

we may write the curl of f as

$$\text{curl}f = \frac{1}{\sqrt{2}}(\text{grad}f - (\text{grad}f)^T).$$

In combination with (21), the above relation yields the usual definition of the curl in 3-D.

The adjoint of the curl, which maps bivector fields to vector fields, is given by the following expression:

$$(\text{curl}^*f)^\wedge = \frac{1}{\sqrt{2}}(\hat{f} - \hat{f}^T)i\xi$$

(note that in the former equation f is a vector field, whereas in the latter, it denotes a tensor).

Finally, we note that in 3-D, under an orthogonal transformation of the domain by $\omega \in \Omega$, the bivector and pseudovector representations [the sides of (21)] transform, respectively, as

$$\frac{1}{\sqrt{2}}\omega^T \begin{bmatrix} 0 & -f_3(\omega^T \cdot) & f_2(\omega^T \cdot) \\ f_3(\omega^T \cdot) & 0 & -f_1(\omega^T \cdot) \\ -f_2(\omega^T \cdot) & f_1(\omega^T \cdot) & 0 \end{bmatrix} \omega$$

$$(\det \omega) \omega \begin{bmatrix} f_1(\omega^T \cdot) \\ f_2(\omega^T \cdot) \\ f_3(\omega^T \cdot) \end{bmatrix}$$

(the determinant captures the sign flip of pseudovectors under parity transformations). The first of these two defines the general law of action of Ω on bivectors in d dimensions

$$[\omega]_b f = \omega^T f(\omega^T \cdot) \omega.$$

B. Curl-like and Divergence-Like Operators and Their Adjoints

Next, we shall give categorized definitions of d -dimensional curl- and divergence-like families of operators and make the connection between these operators and the scalar and vector Laplacians of the previous sections. However, before this, let us first briefly recall, in a single place, the law of action of transformation groups on scalars, vectors, and bivectors. For orthogonal transformation by $\omega \in \Omega$, we have

$$[\omega]_s f = f(\omega^T \cdot) \quad (22)$$

$$[\omega]_v f = \omega f(\omega^T \cdot) \quad (23)$$

$$[\omega]_b f = \omega^T f(\omega^T \cdot) \omega. \quad (24)$$

Note that in the first equation, f is scalar, in the second it is vector, and finally, in the last equation, its values are $d \times d$ antisymmetric matrices.

The actions of T and Σ on the three categories (scalar, vector, and bivector) remain the same in all cases ($f \mapsto f(\cdot - \tau)$ for the former and $f \mapsto f(\sigma^{-1} \cdot)$ for the latter, for all $\tau \in T$ and $\sigma \in \Sigma$).

As noted, we may, in more generality than the previous sections, study the two families of curl- and divergence-like operators and their adjoints. Operators in the former category go from d coordinates to 1, and back by their adjoint; whereas those in the latter go from d coordinates to $d(d-1)/2$ independent coordinates (forming a $d \times d$ antisymmetric matrix) and back to d :

- 1) *Divergences* R_{div} and their adjoints R_{div}^* : These consist of maps from vector to scalar fields and vice versa. In the first case, the invariance equation takes the following form:

$$R_{\text{div}}[S]_v = [S]_s R_{\text{div}};$$

and in the second case, we require

$$R_{\text{div}}^*[S]_s = [S]_v R_{\text{div}}^*.$$

Given our focus on linear regularization operators and the shift-invariance assumption, we can restate the above properties as conditions on the Fourier multipliers of R_{div} and R_{div}^* . Scale invariance in all cases leads to the same equation as (17). With regard to reflection invariance, in place of (18), we have

$$\hat{R}_{\text{div}}(\omega \cdot) = \hat{R}_{\text{div}}(\cdot)\omega^T$$

for the divergence-like operators and

$$\hat{R}_{\text{div}}^*(\omega \cdot) = \omega \hat{R}_{\text{div}}^*(\cdot)$$

for their adjoints. These follow from (22) and (23).

- 2) *Curls* R_{curl} and their adjoints R_{curl}^* : Curl-like operators map vector fields to fields of bivectors. Accordingly, their adjoints map bivectors back to vectors. For the two, we respectively have

$$\begin{aligned} R_{\text{curl}}[S]_v &= [S]_b R_{\text{curl}}; \\ R_{\text{curl}}^*[S]_b &= [S]_v R_{\text{curl}}^*. \end{aligned}$$

In the Fourier multipliers, scale invariance is again reflected by (17). For orthogonal invariance, the equivalents of the preceding pair of equations are respectively

$$\begin{aligned} \left[\hat{R}_{\text{curl}}(\omega \cdot) \right]_{\alpha\beta\gamma} &= \omega_{\alpha i}^T \omega_{j\beta} \omega_{k\gamma}^T \left[\hat{R}_{\text{curl}}(\cdot) \right]_{ijk} \\ \left[\hat{R}_{\text{curl}}^*(\omega \cdot) \right]_{\alpha\beta\gamma} &= \omega_{i\alpha} \omega_{\beta j}^T \omega_{\gamma k} \left[\hat{R}_{\text{curl}}^*(\cdot) \right]_{ijk}. \end{aligned}$$

These are consequences of (22) and (24) (we are using here a light form of Einstein's summation convention, whence repeated indices are summed upon; for instance, $c_{ij} = a_{ik} b_{kj}$ is the product of matrices a_{ij} and b_{ij}). Notice that \hat{R}_{curl} and \hat{R}_{curl}^* are third-rank tensors (linear maps between vectors and matrices) respectively acting on vectors and matrices by

$$[R_{\text{curl}} f]_{ij}^\wedge = [\hat{R}_{\text{curl}}]_{ijk} \hat{f}_k \quad \text{and} \quad [R_{\text{curl}}^* f]_k^\wedge = [\hat{R}_{\text{curl}}^*]_{ijk} \hat{f}_{ij}.$$

Example 1: Fractional Divergences and Gradients: These are denoted by div^γ and grad^γ , respectively, and are defined by the respective symbols given below:

$$|\xi|^\gamma \frac{i\xi^T}{|\xi|} \quad \text{and} \quad |\xi|^\gamma \frac{i\xi}{|\xi|}.$$

The Fractional divergences act on vector fields, mapping them to scalars; gradients do the opposite, with div^γ and $\text{div}^{\gamma*} = -\text{grad}^\gamma$ forming and adjoint pair. The fractional gradient of order 0 (grad^0) is also known as the Riesz transform [22]. \square

Example 2: Fractional Curls and Their Adjoint: We shall denote the fractional curl and its adjoint by curl^γ and $\text{curl}^{\gamma*}$, respectively. They are defined in the Fourier domain according to

$$\begin{aligned} (\text{curl}^\gamma f)^\wedge &= \frac{1}{\sqrt{2}} |\xi|^\gamma \left(\frac{i\xi}{|\xi|} \hat{f}^T - \hat{f} \frac{i\xi^T}{\xi} \right) \\ (\text{curl}^{\gamma*} f)^\wedge &= \frac{1}{\sqrt{2}} |\xi|^\gamma (\hat{f} - \hat{f}^T) \frac{i\xi}{|\xi|}. \end{aligned}$$

These definitions are valid in any number of dimensions ≥ 2 (they are trivial in one dimension). Fractional curls map d -dimensional vectors to $d \times d$ antisymmetric bivectors; adjoint curls go in the opposite direction. \square

One readily verifies that the above examples satisfy the invariances outlined in 1) and 2).

Our claim has been that the considerations of this section are more general than those of the previous two; and yet, until this point, they seem to have been limited to operators mapping vectors to nonvectors and vice-versa. We now show that the former families of Section II (scalar to scalar) and Section IV (vector to vector) can be decomposed in terms of fractional curls and divergences and their adjoints. Specifically, for the scalar fractional Laplacian, we have

$$\begin{aligned} (-\Delta)^\gamma &= \text{div}^\gamma (\text{div}^\gamma)^* \\ &= \text{div}^\gamma (-\text{grad}^\gamma) \end{aligned}$$

and for the generalized fractional vector Laplacian of Section IV, we have

$$\begin{aligned} (-\Delta)_{(c_1, c_2)}^\gamma &= c_1 (\text{div}^\gamma)^* \text{div}^\gamma + c_2 (\text{curl}^\gamma)^* \text{curl}^\gamma \\ &= c_1 (-\text{grad}^\gamma) \text{div}^\gamma + c_2 (\text{curl}^\gamma)^* \text{curl}^\gamma \end{aligned}$$

or, what is the same,

$$\begin{aligned} (\text{div}^\gamma)^* \text{div}^\gamma &= (\text{Id} - \mathbb{P})(-\Delta)^\gamma \\ (\text{curl}^\gamma)^* \text{curl}^\gamma &= \mathbb{P}(-\Delta)^\gamma. \end{aligned}$$

In addition, we record the following factorization results that relate the fractional curls and divergences to combinations of integer-order operators and fractional vector Laplacian $(-\Delta)^\gamma$:

$$\begin{aligned} \text{div}^\gamma &= \text{div}(-\Delta)^{\gamma/2} \\ \text{grad}^\gamma &= (-\Delta)^{\gamma/2} \text{grad} \\ \text{curl}^\gamma &= \text{curl}(-\Delta)^{\gamma/2} \\ \text{curl}^{\gamma*} &= (-\Delta)^{\gamma/2} \text{curl}^*. \end{aligned}$$

We shall not burden ourselves further by trying to find, in complete generality, the equivalents of Theorems 1 and 2 for the curl- and divergence-like families as the cases covered by the above examples appear to us to be sufficiently versatile for applications.

Note, finally, that in order to form regularization functionals similar to (9), which involve curl-like operators, we shall need to define the equivalent of p -norms on $d \times d$ tensor fields. The matrix L_p norms defined in Appendix A perfectly work for this purpose. It is also easy to see that in the case of antisymmetric matrices, the functional obtained in this ways is equal to the vector L_p norm of the upper diagonal elements of the matrix (in particular, in 3-D, this is effectively the same as the norm applied to vector fields). This means that we may alternatively define the same regularization functional in terms of the vector L_p norm of the \star -map of the curl [cf. (21)].

Given all this, a general vector cost functional with multiple regularizers can be written as

$$\begin{aligned} \mathfrak{J}_{(\lambda_1, \lambda_2)}(f; y) = & \mathfrak{d}(f; y) + \sum_{1 \leq i \leq N} \lambda_i \left\| (-\Delta)_{(c_1, i, c_2, i)}^{\gamma_i/2} f \right\|_{p_i}^{q_i} \\ & + \sum_{1 \leq i \leq N'} \lambda'_i \left\| \text{curl}^{\gamma'_i} f \right\|_{p'_i}^{q'_i} \\ & + \sum_{1 \leq i \leq N''} \lambda''_i \left\| \text{div}^{\gamma''_i} f \right\|_{p''_i}^{q''_i}. \end{aligned}$$

(the three p -norms appearing in the above equation are those defined for vectors, bivectors, and scalars, in that order; cf. (7) and Appendix A).

An illustrative example is

$$\begin{aligned} \mathfrak{J}_{(\lambda_1, \lambda_2)}(f; y) = & \mathfrak{d}(f; y) + \lambda_0 \left\| (-\Delta)^{\gamma/2} f \right\|_p^p \\ & + \lambda_1 \left\| (\text{Id} - \mathbb{P})(-\Delta)^{\gamma/2} f \right\|_p^p \\ & + \lambda_2 \left\| \mathbb{P}(-\Delta)^{\gamma/2} f \right\|_p^p \\ & + \lambda_c \left\| \text{curl}^{\gamma} f \right\|_p^p + \lambda_d \left\| \text{div}^{\gamma} f \right\|_p^p \end{aligned}$$

which incorporates independent regularization of the curl- and divergence-free subspaces [see the definition of \mathbb{P} after (20)], as well as the fractional curl and divergence terms. Note that some of λ_i 's may be zero.

We conclude this section by the observation that, as we also saw in (14), in the quadratic case ($p = 2$), the aforementioned functional reduces to the one given at the end of Section IV since, in this particular case, the norm is associated with an inner product, thus allowing us to equate $\|\text{curl}^{\gamma} f\|_2$ with $\|\mathbb{P}(-\Delta)^{\gamma/2} f\|_2$ and $\|\text{div}^{\gamma} f\|_2$ with $\|(\text{Id} - \mathbb{P})(-\Delta)^{\gamma/2} f\|_2$, as can be readily verified using Parseval's identity. This is generally not true for other values of p (but it would have been, had we considered the L_p norms in the Fourier domain in place of the usual spatial L_p norms).

VI. ILLUSTRATION

For the purpose of illustration, we now consider the problem of reconstructing a vector field from noisy measurements—primarily in 3-D ($d = 3$) but also in 2-D—using a quadratic fidelity criterion (consistent with a white Gaussian noise assumption). We shall focus on divergence–curl regularization with different (L_2 versus L_1) norms. We note in passing that in practical problems, higher order regularization, such as the physically motivated second-order divergence–curl regularization of Suter [1], can be of interest, particularly in the context of motion estimation. Here, our primary motivation is to demonstrate and compare the use of L_2 versus L_1 norms, in line with the similar comparison of quadratic versus total-variation-type regularization of scalars that has frequently been made in image processing literature. For this reason, we shall limit ourselves to first-order

differential regularization operators. Specifically, we shall consider the following cost functions:

$$\begin{aligned} \mathfrak{J}^{(p)}(f; Y) = & \sum_n |f(n) - Y[n]|^2 \\ & + \lambda_c \|\text{curl} f\|_p^p + \lambda_d \|\text{div} f\|_p^p \end{aligned} \quad (25)$$

with $p = 1, 2$, where $Y[n]$ s are the measurements (in this section, uppercase letters will be used to denote discrete quantities such as $Y = Y[n] = (Y_1[n], Y_2[n], Y_3[n])$, for n in some subset of \mathbb{Z}^d). In interpreting the above formula when the number of samples and/or estimated values $f(x)$ goes to infinity, some form of normalization or limit argument may become necessary. However, in practice, the number of observations $Y[n]$ will be finite.

The norm applied to the curl in the former equation is a matrix L_p norm, as defined in Appendix A, but in 3-D, we may use the \star -map defined in (21) and rewrite it as a vector norm [cf. (7)] as follows:

$$\begin{aligned} \mathfrak{J}^{(p)}(x; Y) = & \sum_i |f(i) - Y[i]|^2 + \lambda_c \int_{\mathbb{R}^d} \left(\sqrt{|\star \text{curl} f|^2} \right)^p \\ & + \lambda_d \int_{\mathbb{R}^d} \left(\sqrt{|\text{div} f|^2} \right)^p. \end{aligned} \quad (26)$$

For $p = 1$, the mixed $L_2 - L_1$ functional proposed above is in the spirit of TV regularization. It is of interest to compare it against its purely quadratic counterpart, if only to see whether the relative advantage of TV regularization to quadratic regularization in 2-D image denoising carries over to the vector setting.

In three dimensions, the explicit definitions of the curl and the divergence are

$$\begin{aligned} \text{div} f = & \sum_{1 \leq k \leq n} \partial_k f_k = \partial_1 f_1 + \partial_2 f_2 + \partial_3 f_3; \\ \star \text{curl} f = & [\partial_3 f_2 - \partial_2 f_3; \partial_1 f_3 - \partial_3 f_1; \partial_2 f_1 - \partial_1 f_2]. \end{aligned}$$

While guided by the previous continuous formulation, our implementation on a digital computer is necessarily discrete. Although there is room for more sophistication, we shall discretize simply by taking finite differences in place of derivatives while emphasizing that, in practice, the discretization scheme used can play an important role in the numerical solution of inverse problems. It is therefore advisable, in real-world problems, to look at alternatives such as discrete orthogonal decompositions (see [23]).

Let $F_j = F_j[n]$, $j = 1, 2, \dots, d$, denote the reconstruction corresponding to samples of f_j over some discrete domain $\subset \mathbb{Z}^d$. Furthermore, let us denote by $\delta_i F_j$ the finite difference associated with the partial derivative $\partial_i f_j$, that is

$$\delta_i : F_j \mapsto F_j - F_j[\cdot - \hat{e}_i]$$

except at the boundaries where some preferred type of boundary conditions is applied (\hat{e}_i is the i th standard unit vector in \mathbb{R}^d). For future reference, also note the adjoint of δ_i as

$$\delta_i^* : F_j \mapsto F_j - F_j[\cdot + \hat{e}_i]$$

(with appropriate adjoint boundary conditions).

Discrete divergence and curl can be defined in 3-D by the following identities:

$$\begin{aligned}\operatorname{div}_\delta F &= \delta_1 F_1 + \delta_2 F_2 + \delta_3 F_3; \\ \operatorname{curl}_\delta F &= [\delta_3 F_2 - \delta_2 F_3; \delta_1 F_3 - \delta_3 F_1; \delta_2 F_1 - \delta_1 F_2].\end{aligned}$$

The pointwise squared amplitudes of the curl and the divergence that appear under the square root sign in (26) are then discretized as

$$\begin{aligned}|\operatorname{curl}_\delta F[m]|^2 &= \sum_{1 \leq i < j \leq 3} (\delta_i F_j[m] - \delta_j F_i[m])^2 \\ |\operatorname{div}_\delta F[m]|^2 &= \sum_{1 \leq i, j \leq 3} \delta_i F_i[m] \delta_j F_j[m].\end{aligned}$$

Our discrete cost function can be then written as

$$\begin{aligned}\mathfrak{J}_\delta^{(p)}(F; Y) &= \sum_m |F[m] - Y[m]|^2 \\ &+ \lambda_c \sum_m \left(\sqrt{|\operatorname{curl}_\delta F[m]|^2} \right)^p \\ &+ \lambda_d \sum_m \left(\sqrt{|\operatorname{div}_\delta F[m]|^2} \right)^p\end{aligned}\quad (27)$$

(recall that $F[m]$ and $Y[m]$ are vectors, and $|\cdot|$ denotes the Euclidean length; the index m runs over the sampling/reconstruction grid in \mathbb{Z}^d).

For $p = 2$, the problem is quadratic and can be efficiently solved using iterative linear methods. For the L_1 problem, following Figueiredo *et al.* [15], we shall now propose an iterative reweighted least squares (IRLS) approach belonging to the family of majorize–minimize algorithms.

Given some F' with $|RF'| > 0$, the L_1 terms of the functional can be upper bounded as

$$\begin{aligned}\sum \sqrt{|RF|^2} &\leq \sum \sqrt{|RF'|^2} \\ &+ \sum (|RF|^2 - |RF'|^2) / 2\sqrt{|RF'|^2}\end{aligned}\quad (28)$$

(this follows from the inequality $\sqrt{a} \leq \sqrt{a'} + (1/2)(a - a')/\sqrt{a'}$). Let sequence $\tilde{F}_{(n)}$ be defined by

$$\tilde{F}_{(n)} = \arg \min_F \Omega_\delta(F, \tilde{F}_{(n-1)}; Y)\quad (29)$$

where

$$\begin{aligned}\Omega_\delta(F, F'; Y) &:= \sum_m \sum_{1 \leq i \leq 3} F_i[m]^2 \\ &- \sum_m \sum_{1 \leq i \leq 3} 2F_i[m]Y_i[m] \\ &+ \lambda_c \sum_m \frac{|\operatorname{curl}_\delta F[m]|^2}{\sqrt{|\operatorname{curl}_\delta F'[m]|^2}} \\ &+ \lambda_d \sum_m \frac{|\operatorname{div}_\delta F[m]|^2}{\sqrt{|\operatorname{div}_\delta F'[m]|^2}} + \mathfrak{K}_\delta(F'; Y)\end{aligned}$$

TABLE II

COMPARISON OF DENOISING ALGORITHMS IN 3-D; ALGORITHM PARAMETERS WERE OPTIMIZED FOR THE BEST SNR FOR EACH REGULARIZER AND INPUT SNR. (A) GRADIENT FIELD. (B) TUBE AND TORUS.

(a)		
input SNR [dB]	SNR improvement [dB]	
angular error [deg.]	angular error (mean \pm stdev) [deg.]	
	L_1	L_2
0	11.70	11.04
(59.12° \pm 39.93°)	(28.61° \pm 31.46°)	(31.84° \pm 33.95°)
10	7.50	7.78
(37.81° \pm 36.74°)	(16.90° \pm 23.05°)	(20.87° \pm 28.31°)
20	4.49	4.89
(20.22° \pm 28.11°)	(10.03° \pm 15.80°)	(12.40° \pm 21.25°)
(b)		
input SNR [dB]	SNR improvement [dB]	
angular error [deg.]	angular error (mean \pm stdev) [deg.]	
	L_1	L_2
0	8.03	6.37
(12.11° \pm 7.29°)	(5.97° \pm 4.04°)	(5.95° \pm 3.93°)
10	7.96	2.55
(3.82° \pm 2.21°)	(2.58° \pm 1.82°)	(3.16° \pm 2.42°)
20	6.67	0.51
(1.21° \pm 0.70°)	(0.99° \pm 0.71°)	(1.25° \pm 0.85°)

is obtained by majorizing (27) using (28); we have collected all terms depending only on Y and F' in the scalar function \mathfrak{K}_δ , which we may discard when solving (29).

Note that $\Omega(F, F'; Y) = \mathfrak{J}_\delta^{(1)}(F; Y)$. Furthermore, we have

$$\begin{aligned}\mathfrak{J}_\delta^{(1)}(\tilde{F}_{(n)}; Y) &\leq \Omega_\delta(\tilde{F}_{(n)}, \tilde{F}_{(n-1)}; Y) \\ &< \Omega_\delta(\tilde{F}_{(n-1)}, \tilde{F}_{(n-1)}; Y) = \mathfrak{J}_\delta^{(1)}(\tilde{F}_{(n-1)}; Y)\end{aligned}$$

which shows that, with increasing n , the $\mathfrak{J}_\delta^{(1)}(\tilde{F}_{(n)})$'s form a decreasing sequence (in the second inequality, we have used the strict convexity of Ω_δ and assumed that $\tilde{F}_{(n)} \neq \tilde{F}_{(n-1)}$).

For fixed F' , the minimizer of $\Omega_\delta(F, F'; Y)$ over F is the solution of the linear system of equations obtained by setting all of the derivatives of Ω_δ equal to zero. To see this, let us first define the following:

$$c'_m = \sqrt{|\operatorname{curl}_\delta F'[m]|^2} \quad d'_m = \sqrt{|\operatorname{div}_\delta F'[m]|^2}.$$

Furthermore, let

$$\begin{aligned}C'_k F[p] &:= \sum_{1 \leq i \leq 3} \delta_i^* \frac{\delta_i F_k[p] - \delta_k F_i[p]}{c'_p} \\ D'_k F[p] &:= \delta_k^* \frac{\sum_{1 \leq i \leq 3} \delta_i F_i[p]}{d'_p}.\end{aligned}$$




Fig. 1. “Gradient” phantom; see text for a description of the experiment. (a) Noisy field (0 dB SNR). (b) Denoised field, using L_1 regularization (11.70 dB SNR).




Fig. 2. “Tube and torus” phantom; see text for a description of the experiment. (a) Noisy field (0 dB SNR). (b) Denoised field, using L_1 regularization (9.01 dB SNR).

After some algebraic simplification, one can write

$$\frac{\partial}{\partial F_k[p]} \mathfrak{Q}_\delta(F, F'; Y) = 2(F_k[p] - Y_k[p]) + 2\lambda_C C'_k F[p] + 2\lambda_D D'_k F[p].$$

The system of equations defined by

$$\frac{\partial}{\partial F_k[p]} \mathfrak{Q}_\delta(F, F'; Y) = 0, \quad \text{for } k = 1, 2, 3 \text{ and all } p \quad (30)$$

thus corresponds to linear system $AF = Y$ (shorthand for $(AF)_k[p] = Y_k[p]$, for all p). This system may then be solved using a variety of methods (conjugate gradient (CG), multi-grid-preconditioned GMRES, etc.). In implementation, one may add a small ϵ to numerators and denominators to avoid division by zero.

To summarize, the complete algorithm for L_1 regularized denoising consists of a number of *outer* cycles in accordance

with (29), which sequentially reduce cost functional \mathfrak{Q}_δ . The n th outer iteration takes measurements Y and the output of the $(n-1)$ th iteration ($\tilde{F}_{(n-1)}$) as inputs and then moves in the direction of minimizing $\mathfrak{Q}_\delta(\cdot, \tilde{F}_{(n-1)}, Y)$. This local minimization corresponds to a linear system, as specified in (30). Within each outer iteration, this system is then (approximately) solved using a number of *inner* iterations of some iterative linear solver.

A. Simulation and Results

We implemented the previously described scheme in MATLAB (The MathWorks, Inc., Natick, MA) in 2- and 3-D. As experiments, we considered the denoising of phantoms corrupted by different levels of white Gaussian noise. λ_C and λ_D were optimized for best mean squared error (MSE) performance. In simulation, the true MSE for a given choice of λ_C and λ_D can be calculated using an oracle. In practice, even though the ground truth is not known and the true MSE is therefore not accessible, so long as the white Gaussian




Fig. 3. Amplitude cross sections: “gradient” phantom, comparing L_1 and L_2 denoising. (a) Original. (b) Noisy (0 dB SNR). (c) L_1 denoised (11.70 dB improvement). (d) L_2 denoised (11.04 dB improvement).




Fig. 4. Amplitude cross sections: “tube and torus” phantom, comparing L_1 and L_2 denoising. (a) Original. (b) Noisy (10 dB SNR). (c) L_1 denoised (7.96 dB improvement). (d) L_2 denoised (2.55 dB improvement).

noise assumption remains valid, a highly accurate estimate of the MSE can be obtained using Monte Carlo techniques that approximate Stein’s unbiased risk estimate, as described in Ramani *et al.* [24] (see also [25]). This estimate comes at the cost of solving an extra denoising problem for each choice of λ_C and λ_D , but in terms of effectiveness in predicting the best values of λ_C and λ_D , we found it to be indistinguishable from the oracle in our experiments.

Results are reported in Table II and in Figs. 1–5 (3-D graphics were generated using ParaView 3.8.0 [26]). The phantoms, and high-resolution images of their noisy and denoised versions, are available online, at the web address <http://bigwww.epfl.ch/tafti/gal/vreg/>.

The first 3-D phantom, presented in Fig. 1, consists of the gradient field of the potential

$$\phi_{3D}(x_1, x_2, x_3) = x_1 x_2 e^{-|x|^2}.$$

The second 3-D phantom, depicted in Fig. 2, models fully developed laminar flow (with a parabolic profile) in a tube, encircled by constant flow inside a torus.

We solved the L_1 version of the denoising problem using the previously described IRLS scheme, with eight external cycles per (29) and 600 CG inner iterations per cycle to solve the linearized problem in each step. The L_2 problem was solved to convergence using CG iterations.

In Figs. 3 and 4, we show the amplitude profile of the L_1 and L_2 reconstructions of the two 3-D phantoms. These reconstructions are also compared in Table II in terms of SNR improvement after denoising (with λ_C and λ_D optimized for the best

SNR performance) and mean angular error. The latter performance measure is defined as the average pointwise angle between the ground truth and the reconstruction (see [27]).

The point we wish to highlight here is that L_1 regularization performs remarkably well for the second phantom, which features discontinuities in the flow, while being almost comparable to L_2 regularization for the first (smooth) phantom. The former regularization also better preserves small details and discontinuities at flow boundaries, which are smoothed in L_2 denoising. On the other hand, not unexpectedly, L_2 denoising produces slightly higher SNRs in the case of the smooth “gradient” phantom, although L_1 regularization is still quite comparable in terms of SNR and even yields smaller angular errors.

As previously hinted, we took advantage of the availability of the ground truth to optimize parameters λ_C and λ_D for the best SNR, for which purpose we used a bracketing search method (it also bears reminding that the parameters were therefore *not* optimized for our second quality criterion, which is the mean angular error). The parameter values obtained in our experiments are tabulated in Table III. We remark that the superior performance of the L_1 algorithm is in spite of the fact that, in contrast to the L_2 case, the experimentally obtained parameters λ_C and λ_D for the L_1 problem may be, in fact, suboptimal, primarily as a consequence of the L_1 problem being typically solved only partially by fixing the number of iterations (computational budget) in advance, meaning that due to the variable state of convergence, SNR performance fluctuates about its optimum, thus breaking the working assumptions of typical optimization algorithms used to optimize λ_C and λ_D . It is also worth noting that terminating the scheme before full conver-

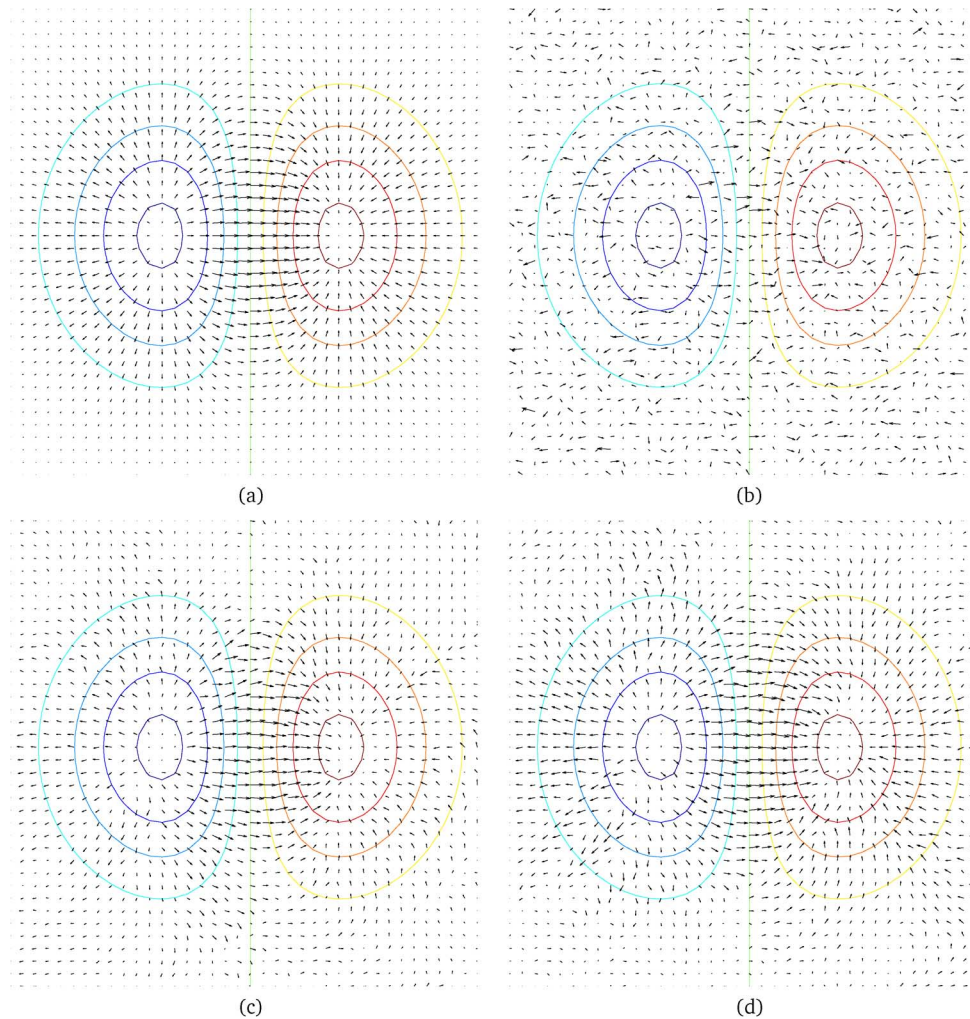


Fig. 5. Denoising in 2-D with L_1 and L_2 regularization applied to the noisy gradient of ϕ_{2D} (see text) with contour lines of ϕ_{2D} overlaid in color. (a) Original. (b) Noisy. (c) L_1 denoising output (12.74 dB SNR improvement). (d) L_2 denoising output (12.58 dB SNR improvement).

TABLE III
OPTIMAL λ_C AND λ_D PAIRS USED TO OBTAIN THE RESULTS IN TABLE II

input SNR [dB]	λ_C, λ_D			
	gradient field		tube and torus	
	L_1	L_2	L_1	L_2
0	0.511, 0.715	2.435, 1.802	0.335, 1.469	0.536, 6.890
10	0.241, 0.473	0.954, 0.677	0.060, 0.047	0.095, 0.858
20	0.009, 0.005	0.384, 0.243	0.022, 0.017	0.012, 0.058

gence can itself be seen as an additional source of regularization; the optimal parameters λ_C and λ_D therefore also depend on the state of convergence of the problem.

As a further demonstration of potential, in Fig. 5, we provide a sample output of 2-D vector field denoising. The phantom used in this case was the gradient of the potential function

$$\phi_{2D}(x_1, x_2) = x_1 e^{-|x|^2}$$

(contour lines of ϕ_{2D} are superimposed in color). For the example shown in Fig. 5, we observed an SNR improvement of

12.74 dB with L_1 regularization, compared with an improvement of 12.58 dB when using quadratic (L_2) regularization. We note that reconstruction of 2-D vector fields can have applications beyond denoising, for instance in image registration and motion estimation, although in the latter case, temporal regularization also needs to be considered.

VII. CONCLUSION

In this paper, we have studied the question of designing regularization functionals for variational reconstruction of vector fields. We approached this problem on the basis of requiring that the regularization functional satisfy certain geometric invariance properties, which we justified from different angles. To set the stage for our derivations, we first addressed some commonalities of invariant regularization in scalar and vector settings—followed by a derivation of the general form of invariant regularizers for scalar fields—before specializing to the problem of invariant vector regularization. The vector regularization functionals that we have derived consist of combinations of (possibly fractional) curl- and divergence-like operators and their adjoints, wrapped in scalar, vector, and/or matrix L_p norms

(also introduced in this paper). The generalized vector Laplacians of [20] also fall within this framework.

Finally, we have presented an application of the proposed framework to the problem of vector field denoising in 2- and 3-D, where we gave a natural generalization of L_2 (quadratic), as well as L_1 (TV-type) regularization for vector fields. While our preliminary results already show a systematic advantage of L_1 over L_2 regularization in the aforementioned problem, many interesting questions about the choice of higher order regularization functionals for data with specific structure remain open. Moreover, in addition to vector denoising, the proposed construction can find applications in a variety of other problems, which we did not study here. Examples include reconstruction of vector fields from nonuniform and incomplete (scalar) measurements, deconvolution, estimation of optical flow, and image registration. Another possible direction for future investigations is the incorporation of temporal regularization in the formulation.

APPENDIX

A. Matrix L_p Norms and Spaces

The vector L_p norms defined in (7) are special cases of the L_p norms for the matrix-valued functions $\mathbb{R}^d \rightarrow \mathbb{R}^{n \times m}$ that we define as follows (ρ denotes the spectral radius):

$$\|f\|_p := \begin{cases} \left(\int_{\mathbb{R}^d} \text{Tr} \left([f(x)^H f(x)]^{p/2} \right) dx \right)^{1/p}; & p < \infty, \\ \|\rho(f^H f)^{1/2}\|_\infty. & p = \infty \end{cases}$$

This definition is motivated by a matrix Young inequality due to [28], i.e.,

$$\text{Tr} \left[(Y^H X^H X Y)^{1/2} \right] \leq \frac{1}{p} \text{Tr} \left[(X^H X)^{p/2} \right] + \frac{1}{p'} \text{Tr} \left[(Y^H Y)^{p'/2} \right]$$

where $1/p + 1/p' = 1$. The preceding inequality can be used to prove a version of Hölder's inequality for L_p spaces of matrix-valued functions. Matrix L_p spaces are then defined in the standard manner. They have similar properties to scalar L_p spaces (completeness, inner product structure for $p = 2$, duality between L_p and $L_{p'}$ with $1/p + 1/p' = 1$ via the bilinear form $\langle f, g \rangle := \int_{\mathbb{R}^d} f^T g$, etc.).

Definitions of matrix ℓ_q norms and spaces are obtained by replacing the integrals with sums.

B. Proof of Theorem 2

By (18) (this is the only place in this proof where we use invariance to improper rotations), we have

$$\hat{R}(-\xi) = (-I)\hat{R}(\xi)(-I) = \hat{R}(\xi). \quad (31)$$

Next, fix $\xi \neq 0$, and let ξ_i^\perp , $1 \leq i \leq d-1$, be $d-1$ pairwise orthogonal vectors in \mathbb{R}^d , all perpendicular to ξ and with $|\xi_i^\perp| = |\xi|$. We define the following rotation matrices:

$$\omega_i = I - 2\xi\xi^T/|\xi|^2 - 2\xi_i^\perp (\xi_i^\perp)^T/|\xi|^2, \quad 1 \leq i \leq d-1.$$

Each ω_i is a simple rotation by 180° in the $\xi \wedge \xi_i^\perp$ plane. In particular, $\omega_i \xi = -\xi$. We also define, for each pair $i \neq j$, the 90° rotation matrix

$$\omega_{ij} = I - \frac{\xi_i^\perp (\xi_i^\perp)^T}{|\xi|^2} - \frac{\xi_j^\perp (\xi_j^\perp)^T}{|\xi|^2} + \frac{\xi_j^\perp (\xi_i^\perp)^T}{|\xi|^2} - \frac{\xi_i^\perp (\xi_j^\perp)^T}{|\xi|^2}.$$

ω_{ij} maps $\xi_i^\perp \mapsto \xi_j^\perp \mapsto -\xi_i^\perp$ and leaves ξ fixed (in this proof, ω_{ij} and ξ_i^\perp denote entire matrices and vectors, respectively, and not the entries of some unspecified matrix ω or vector ξ^\perp).

Note that matrices ω_i , $1 \leq i \leq d-1$, commute pairwise; also, by (18) and (31), we have

$$\hat{R}(\xi)\omega_i = \hat{R}(-\xi)\omega_i = \hat{R}(\omega_i\xi)\omega_i = \omega_i\hat{R}(\xi)$$

which shows that ω_i 's commute with $\hat{R}(\xi)$ as well. Since, for $d > 2$, vectors ξ , $\xi_1^\perp, \dots, \xi_{d-1}^\perp$ are precisely the common eigenvectors of $\omega_1, \dots, \omega_{d-1}$, they must also be eigenvectors of $\hat{R}(\xi)$; in particular, ξ is an eigenvector of $\hat{R}(\xi)$. Denote its corresponding eigenvalue by $\mu_1 = \mu_1(\xi)$. By taking the transpose of (18) and applying the same argument, we can show that ξ is also an eigenvector of $\hat{R}(\xi)^T$. Its corresponding eigenvalue, temporarily denoted by $\mu_1'(\xi)$, is equal to $\mu_1(\xi)$ since

$$\mu_1' \xi^T \xi = \xi^T \hat{R}(\xi) \xi = \mu_1 \xi^T \xi.$$

We similarly denote the eigenvalue of ξ_i^\perp by $\mu_{2,i}(\xi)$.

Alternatively, to find the eigenvectors of $\hat{R}(\xi)$, we might note that $\hat{R}(\xi)$ commutes with all ω_{ij} 's as

$$\hat{R}(\xi)\omega_{ij} = \hat{R}(\omega_{ij}\xi)\omega_{ij} = \omega_{ij}\hat{R}(\xi)$$

and since ξ is an eigenvector of all ω_{ij} 's with eigenvalue 1 (it is their only common eigenvector), $\hat{R}(\xi)\xi$ must be a common eigenvector of all ω_{ij} 's; thus, $\hat{R}(\xi)\xi = \mu_1(\xi)\xi$ for some scalar eigenvalue $\mu_1(\xi)$. Then, to show that the ξ_i^\perp 's are also eigenvectors of $\hat{R}(\xi)$, we observe that

$$\begin{aligned} \hat{R}(\xi)\xi_i^\perp &= \hat{R}(-\xi)\xi_i^\perp = \hat{R}(\omega_i\xi)\xi_i^\perp \\ &= \omega_i\hat{R}(\xi)\omega_i^{-1}\xi_i^\perp = \omega_i\hat{R}(\xi)(-\xi_i^\perp) = -\omega_i\hat{R}(\xi)\xi_i^\perp; \end{aligned}$$

whereby $(I + \omega_i)\hat{R}(\xi)\xi_i^\perp = 0$. This shows that $\hat{R}(\xi)\xi_i^\perp$ lies in the kernel of $I + \omega_i$. However, the kernel of $I + \omega_i$ exactly corresponds to the span of $\{\xi, \xi_i^\perp\}$. We can therefore write $\hat{R}(\xi)\xi_i^\perp = \alpha_i\xi + \mu_{2,i}\xi_i^\perp$ for some α_i and $\mu_{2,i}$, but then, $\xi^T\hat{R}(\xi)\xi_i^\perp = \alpha_i|\xi|^2$; we also have $\xi^T\hat{R}(\xi)\xi_i^\perp = \mu_1\xi^T\xi_i^\perp = 0$. The last two equations show that $\alpha_i = 0$, that is, we have $\hat{R}(\xi)\xi_i^\perp = \mu_{2,i}(\xi)\xi_i^\perp$. $\hat{R}(\xi)$ therefore has $\xi, \xi_1^\perp, \dots, \xi_{d-1}^\perp$ as d eigenvectors with respective eigenvalues of $\mu_1(\xi), \mu_{2,1}(\xi), \dots, \mu_{2,d-1}(\xi)$.

Next, we show that all $\mu_{2,i}$'s are equal to some $\mu_2 = \mu_2(\xi)$: by (18), we have

$$\begin{aligned} \mu_{2,j}\xi_j^\perp &= \hat{R}(\xi)\xi_j^\perp = \hat{R}(\omega_{ij}\xi)\xi_j^\perp = \omega_{ij}\hat{R}(\xi)\omega_{ij}^{-1}\xi_j^\perp \\ &= \omega_{ij}\hat{R}(\xi)\xi_i^\perp = \mu_{2,i}\omega_{ij}\xi_i^\perp = \mu_{2,i}\xi_j^\perp \end{aligned}$$

proving that all $\mu_{2,i}(\xi)$'s are equal as claimed. Putting everything together, we find that $\hat{R}(\xi)$ has the orthogonal eigenvectors

$\xi, \xi_1^\perp, \dots, \xi_{d-1}^\perp$ with eigenvalues $\mu_1(\xi)$ for ξ and $\mu_2(\xi)$ for the remaining vectors. Its eigendecomposition is, therefore, of the form given in (19) (for $d = 2$, we can make a similar demonstration of the above decomposition by working with the reflection matrix with axis ξ , instead of ω_{ij} 's). Finally, for (17) and (18) to hold, μ_1 and μ_2 must be rotation invariant and homogeneous of degree γ . Thus, by Theorem 1, we have $\mu_i(\xi) = c_i|\xi|^\gamma$ for some c_i and $i = 1, 2$. ■

REFERENCES

- [1] D. Suter, "Motion estimation and vector splines," in *Proc. IEEE CVPR*, Jun. 1994, pp. 939–942.
- [2] M. Arigovindan, M. Sühling, C. Jansen, P. Hunziker, and M. Unser, "Full motion and flow field recovery from echo Doppler data," *IEEE Trans. Med. Imag.*, vol. 26, no. 1, pp. 31–45, Jan. 2007.
- [3] A. F. Stalder, M. F. Russe, A. Frydrychowicz, J. Bock, J. Hennig, and M. Markl, "Quantitative 2D and 3D phase contrast MRI: Optimized analysis of blood flow and vessel wall parameters," *Magn. Reson. Med.*, vol. 60, no. 5, pp. 1218–1231, Nov. 2008.
- [4] T. Poggio, V. Torre, and C. Koch, "Computational vision and regularization theory," *Nature*, vol. 317, no. 6035, pp. 314–319, Sep./Oct. 1985.
- [5] L. Alvarez, R. Deriche, T. Papadopoulos, and J. Sánchez, "Symmetrical dense optical flow estimation with occlusions detection," *Int. J. Comput. Vis.*, vol. 75, no. 3, pp. 371–385, Dec. 2007.
- [6] G. Christensen, R. Rabbitt, and M. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. Image Process.*, vol. 5, no. 10, pp. 1435–1447, Oct. 1996.
- [7] C. Sánchez Sorzano, P. Thévenaz, and M. Unser, "Elastic registration of biological images using vector-spline regularization," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 4, pp. 652–663, Apr. 2005.
- [8] J. Duchon, "Splines minimizing rotation-invariant semi-norms in Sobolev spaces," in *Constructive Theory of Functions of Several Variables*, A. Dold and B. Eckmann, Eds. New York: Springer-Verlag, 1977, pp. 85–100.
- [9] D. Suter and F. Chen, "Left ventricular motion reconstruction based on elastic vector splines," *IEEE Trans. Med. Imag.*, vol. 19, no. 4, pp. 295–305, Apr. 2000.
- [10] M. Arigovindan, "Variational reconstruction of vector and scalar images from non-uniform samples," Ph.D. dissertation, EPFL, Biomedical Imaging Group, Lausanne, Switzerland, 2005, <http://library.epfl.ch/theses/?nr=3329>.
- [11] F. Dodu and C. Rabut, "Irrotational or divergence-free interpolation," *Numer. Math.*, vol. 98, no. 3, pp. 477–498, Sep. 2004.
- [12] L. Amodè and M. N. Benbourhim, "A vector spline approximation," *J. Approx. Theory*, vol. 67, no. 1, pp. 51–79, Oct. 1991.
- [13] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenom.*, vol. 60, no. 1–4, pp. 259–268, Nov. 1992.
- [14] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 3, pp. 367–383, Mar. 1992.
- [15] M. A. T. Figueiredo, J. B. Dias, J. P. Oliveira, and R. D. Nowak, "On total variation denoising: A new majorization-minimization algorithm and an experimental comparison with wavelet denoising," in *Proc. IEEE Int. Conf. Image Process.*, Atlanta, GA, Oct. 2006, pp. 2633–2636.
- [16] M. A. T. Figueiredo, J. M. Bioucas-Dias, and R. D. Nowak, "Majorization-minimization algorithms for wavelet-based image restoration," *IEEE Trans. Image Process.*, vol. 16, no. 12, pp. 2980–2991, Dec. 2007.
- [17] U. Grenander, *Abstract Inference*. New York: Wiley, 1981.
- [18] I. M. Gel'fand and G. E. Shilov, *Generalized Functions*. New York: Academic, 1964.
- [19] J.-Y. Na and S.-Y. Chung, "Homogeneous generalized functions which are rotation invariant," *J. Math. Kyoto Univ.*, vol. 40, no. 1, pp. 155–163, 2000.
- [20] P. D. Tafti and M. Unser, "Fractional Brownian vector fields," *SIAM Multisc. Model. Simul.*, vol. 8, no. 5, pp. 1645–1670, 2010.
- [21] A. W. McDavid and C. D. McMullen, "Generalizing Cross Products and Maxwell's Equations to Universal Extra Dimensions [Online]. Available: <http://arxiv.org/abs/hep-ph/0609260>
- [22] M. Unser and D. Van De Ville, "Wavelet steerability and the higher-order Riesz transform," *IEEE Trans. Image Process.*, vol. 19, no. 3, pp. 636–652, Mar. 2010.
- [23] J. Yuan, C. Schnörr, and E. Mëmin, "Discrete orthogonal decomposition and variational fluid flow estimation," *J. Math. Imag. Vis.*, vol. 28, no. 1, pp. 67–80, May 2007.
- [24] S. Ramani, T. Blu, and M. Unser, "Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Trans. Image Process.*, vol. 17, no. 9, pp. 1540–1554, Sep. 2008.
- [25] D. A. Girard, "The fast Monte-Carlo cross-validation and C_L procedures: Comments, new results and application to image recovery problems," *Comput. Stat.*, vol. 10, no. 3, pp. 205–258, 1995.
- [26] A. H. Squillacote, *The ParaView Guide: A Parallel Visualization Application*, 3rd ed. Clifton Park, NY: Kitware, Inc., 2008.
- [27] J. L. Barron, D. J. Fleet, and S. S. Beauchemin, "Performance of optical flow techniques," *Int. J. Comput. Vis.*, vol. 12, no. 1, pp. 43–77, Feb. 1994.
- [28] T. Ando, "Matrix Young inequalities," in *Operator Theory in Function Spaces and Banach Lattices*. Basel, Switzerland: Birkhäuser, 1995, pp. 33–38.



Pouya Dehghani Tafti (M'00) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2003 and the M.A.Sc. (Distinction) degree in electrical engineering from McMaster University, Hamilton, ON, Canada, in 2006. He is currently working toward the Ph.D. degree in the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.

Since joining the Biomedical Imaging Group in 2006, he has worked on stochastic and variational models based on geometrical invariances, their connections with wavelets and splines, and their applications in image and vector field reconstruction and (bio)medical imaging.



Michael Unser (M'89–SM'94–F'99) received the M.S. (*summa cum laude*) and Ph.D. degrees in electrical engineering from the École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 1981 and 1984, respectively.

From 1985 to 1997, he worked as a Scientist with the National Institutes of Health, Bethesda, MD. He is currently a Full Professor and the Director of the Biomedical Imaging Group, EPFL. His main research area is biomedical image processing. He has a strong interest in sampling theories, multiresolution algorithms, wavelets, and the use of splines for image processing. He has published over 150 journal papers on those topics.

Dr. Unser is currently a member of the editorial boards of the Foundations and Trends in Signal Processing and the Society for Industrial and Applied Mathematics *Journals on Imaging Sciences* and *Sampling Theory in Signal and Image Processing*. He coorganized the first IEEE International Symposium on Biomedical Imaging in 2002. He was the founding chair of the technical committee of the IEEE Signal Processing Society on Bio Imaging and Signal Processing. He has held the position of an Associate Editor-in-Chief of the IEEE TRANSACTIONS ON MEDICAL IMAGING from 2003 to 2005 and has served as an Associate Editor of the same journal from 1999 to 2002 and from 2006 to 2007, the IEEE TRANSACTIONS ON IMAGE PROCESSING from 1992 to 1995, and the IEEE Signal Processing Letters from 1994 to 1998. He was the recipient of the 1995 and 2003 Best Paper Awards and the 2000 Magazine Award from the IEEE Signal Processing Society and is one of the Institute for Scientific Information's Highly Cited authors in Engineering (<http://isihighlycited.com>).